Centre de recherche sur les transports (C.R.T.)
Centre for Research on Transportation

CRT-2004-20

A TABU SEARCH HEURISTIC FOR THE VEHICLE ROUTING PROBLEM WITH TWO-DIMENSIONAL LOADING CONSTRAINTS

by

Michel Gendreau
Manuel Iori
Gilbert Laporte
Silvano Martello

September 2004

Université de Montréal
C.P. 6128, succursale Centre-ville, Montréal QC H3C 3J7 Canada
Téléphone : (514) 343-7575 / Télécopieur : (514) 343-7121

crt@crt.umontreal.ca
www.crt.umontreal.ca

HEC ÉCOLE POLYTECHNIQUE MONTRÉAL Université de Montréal

PUBLICATION
CRT-2004-20

A TABU SEARCH HEURISTIC FOR THE VEHICLE ROUTING PROBLEM WITH TWO-DIMENSIONAL LOADING CONSTRAINTS

by

Michel Gendreau[1]
Manuel Iori[2]
Gilbert Laporte[3]
Silvano Martello[2]

The research was supported in part by the Ministero dell'Istruzione, delle'Università e della Ricerca (MIUR), the Consiglio Nazionale delle Ricerche (CNR) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

1. Département d'informatique et de recherche opérationnelle and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7 (michelg@crt.umontreal.ca)

2. Dipartimento di Elettronica, Informatica e Sistemistica, Università degli Studi di Bologna, Viale Risorgimento 2, 40136 Bologna, Italy (miori@deis.unibo.it, smartello@deis.unibo.it)

3. Canada Research Chair in Distribution Management, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal, Canada H3T 2A7, and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7 (gilbert@crt.umontreal.ca)

Dépôt légal – Bibliothèque nationale du Québec
Bibliothèque nationale du Canada, 2004

105-00424/

Abstract

The article addresses the well-known Capacited Vehicle Routing Problem (CVRP), in the special case where the demand of a customer consists of a certain number of two-dimensional weighted items. The problem calls for the minimization of the cost of transportation needed for the delivery of the goods demanded by the customers, and carried out by a fleet of vehicles based at a central depot. In order to accommodate all items on the vehicles, a feasibility check of the two-dimensional packing (2L) must be executed on each vehicle. The overall problem, denoted as 2L-CVRP, is NP-hard and particularly difficult to solve in practice. We propose a Tabu Search heuristic, in which the WL-CVRP is solved through heuristics, lower bounds and a truncated branch-and-bound procedure. The effectiveness of the algorithm is demonstrated through extensive computational experiments.

Keywords: Vehicle routing problem, two-dimensional packing, tabu search heuristic.

Résumé

Cet article porte sur le problème bien connu de tournées de véhicules (PTV) dans le cas particulier où la demande des clients comprend des items à deux dimensions. Dans ce problème, on cherche à minimiser le coût de transport requis pour la livraison des demandes aux clients et effectué par une flotte de véhicules basés à un dépôt central. Afin de placer tous les items dans les véhicules, on doit soumettre chaque solution potentielle à un test d'admissibilité (2D) pour chaque véhicule. Le problème global, dénoté par 2D-PVC, est NP-difficile et particulièrement ardu à résoudre en pratique. L'article propose une méthode de recherche avec tabous dans lequel le 2D-PVC est résolu à l'aide d'heuristique, de bornes inférieures et d'une procédure de branch-and-bound. L'efficacité de l'algorithme est démontrée à partir d'expérience de calcul sur des problèmes tests.

Mots-clés : problème de tournées de véhicules, problème de mise en boîtes à deux dimensions, heuristique de recherche avec tabous.

1 Introduction

One of the most frequently studied combinatorial optimization problems is the *Capacitated Vehicle Routing Problem* (CVRP), which calls for the optimization of the delivery of goods, demanded by a set of clients, and operated by a fleet of vehicles of limited capacity based at a central depot. The application of this model to real world problems is limited by the presence of many additional constraints. In particular, in the CVRP client demands are expressed by an integer value, representing the total weight of the items to be delivered, while in real-world instances demands consist of lots of items characterized both by a weight and a three-dimensional volume.

The problem of loading or unloading these three-dimensional general-shape lots to and from vehicles is not trivial and can lead to impractical models for which it is impossible to design effective algorithms. It is, however, possible to simplify the problem. First, it is quite usual to place items into rectangular-shaped boxes. Additionally, in many cases boxes (e.g., pallets) cannot be stocked one on top of the other because of their fragility or weight. We thus restrict our study to the loading of two-dimensional, weighted and rectangular-shaped items.

The resulting problem, which combines the feasible loading or unloading of such items into vehicles and the minimization of transportation costs, is particularly relevant for freight distribution companies. From an algorithmic point of view, the combination of these two areas of combinatorial optimization (vehicle routing and two-dimensional packing) is challenging. We consider two variants of the problem (see Section 2), one of which was first addressed by Iori, Salazar González and Vigo [20]. Following the notation of these authors, the problem will be denoted here as 2L-CVRP (*Two-Dimensional Loading Capacitated Vehicle Routing Problem*).

The vehicle routing problem has been deeply investigated since the seminal work of Dantzig and Ramser [9]. Branch-and-cut algorithms are nowadays the best exact approaches (see, e.g., Toth and Vigo [27]), but are able to systematically solve to optimality only relatively small instances. The problem of loading two-dimensional items into the vehicle is closely related to various packing problems. In particular:

- the *Two-Dimensional Bin Packing Problem* (2BPP): Pack a given set of rectangular items into the minimum number of identical rectangular bins. Exact approaches for the 2BPP are generally based on branch-and-bound techniques and are able to solve instances with up to 100 items (but some instances with 20 items remain unsolved). Recent exact algorithms and lower bounds have been proposed by Martello and Vigo [24], Fekete and Schepers [11], Boschetti and Mingozzi [1, 2] and Pisinger and Sigurd [25].

- the *Two-Dimensional Strip Packing Problem* (2SP): Pack a given set of rectangular items into a strip of given width and infinite height so as to minimize the overall height of the packing. An exact algorithm proposed by Martello, Monaci and Vigo [22] can solve instances with up to 200 items, although some instances with 30 items are still unsolved.

Both the CVRP and the 2BPP are strongly NP-hard and very difficult to solve in practice. The same holds for the 2L-CVRP. To our knowledge, the only available solution methodology for the 2L-CVRP is the exact approach developed by Iori, Salazar González and Vigo [20]: it consists of a branch-and-cut algorithm that handles the routing aspects of the problem, combined with a nested branch-and-bound procedure to check the feasibility of the loadings. The algorithm can solve instances involving up to 30 clients and 90 items. Since these limits are not reasonable for real-life problems, it is natural to consider heuristic and metaheuristic techniques. In this paper we propose a Tabu Search algorithm for the 2L-CVRP.

Several metaheuristic approaches are available for the CVRP (see Cordeau et al. [5] for a recent survey): Tabu Search, Genetic Algorithms, Simulated and Deterministic Annealing, Ant Colonies and Neural Networks. Very good results were obtained through Tabu Search (Cordeau and Laporte [7]). Metaheuristic techniques have also been widely applied to packing problems. The 2BPP was solved through Simulated Annealing by Dowsland [10] and through Tabu Search by Lodi, Martello and Vigo [21]. Tabu Search and Genetic Algorithms for the 2SP were developed by Iori, Martello and Monaci [19].

The success of Tabu Search both for routing and packing problems influenced our choice in favor of such technique for the 2L-CVRP. An additional consideration is that even the problem of finding a feasible solution to the 2L-CVRP (i.e., a partition of the clients into subsets that satisfy the weight capacity of the vehicles and for which a feasible two-dimensional loading exists) is already strongly NP-hard and very difficult to solve in practice. Techniques based on populations of solutions, such as Genetic Algorithms and Scatter Search, would be forced to deal with a high number of infeasible solutions, and hence do not seem very promising.

In Section 2 we introduce our notation and give a detailed description of the addressed problems. In Section 3 we present the proposed Tabu Search approach. The intensification phase, which constitutes a very important component of the algorithm, is described in detail in Section 4. Computational results on small-size and large-size instances are given in Section 5.

2 Problem Description

We are given a complete undirected graph $G = (V, E)$, where V is a set of $n + 1$ vertices corresponding to the depot (vertex 0) and to the clients (vertices $1, \ldots, n$), and E is the set of edges (i, j), each having an associated cost c_{ij}. There are v identical vehicles, each having a weight capacity D and a rectangular loading surface of width W and height H. Let $A = W \times H$ denote the loading area. The demand of client i ($i = 1, \ldots, n$) consists of m_i items of total weight d_i: item $I_{i\ell}$ ($\ell = 1, \ldots, m_i$) has width $w_{i\ell}$ and height $h_{i\ell}$. Let $a_i = \sum_{\ell=1}^{m_i} w_{i\ell} h_{i\ell}$ denote the total area of the lot demanded by client i.

We assume that the items have a fixed orientation, i.e., they must be packed with the w-edge (resp. the h-edge) parallel to the W-edge (resp. the H-edge) of the loading surface. This restriction is quite usual in logistics when pallet loading is considered, since

there can be a fixed orientation for the entering of the forks that prohibits rotations during the loading and unloading operations. We also assume, as is usual for the VRP, that each client must be served by a single vehicle. When a vehicle k is assigned a tour that includes a client set $S(k) \subseteq \{1, \ldots, n\}$, the two following constraints must be satisfied:

Weight constraint: the total weight $\sum_{i \in S(k)} d_i$ must not exceed the vehicle capacity D;

Loading constraint: there must be a feasible (non-overlapping) loading of the items in

$$I(k) = \bigcup_{i \in S(k)} \bigcup_{\ell \in \{1,\ldots,m_i\}} I_{i\ell} \tag{1}$$

into the $W \times H$ loading area.

The objective is to find a partition of the clients into at most v subsets and, for each subset $S(k)$, a route starting and ending at the depot (vertex 0) such that both conditions above hold, and the total cost of the edges in the routes is a minimum. An example is depicted in Figure 1.

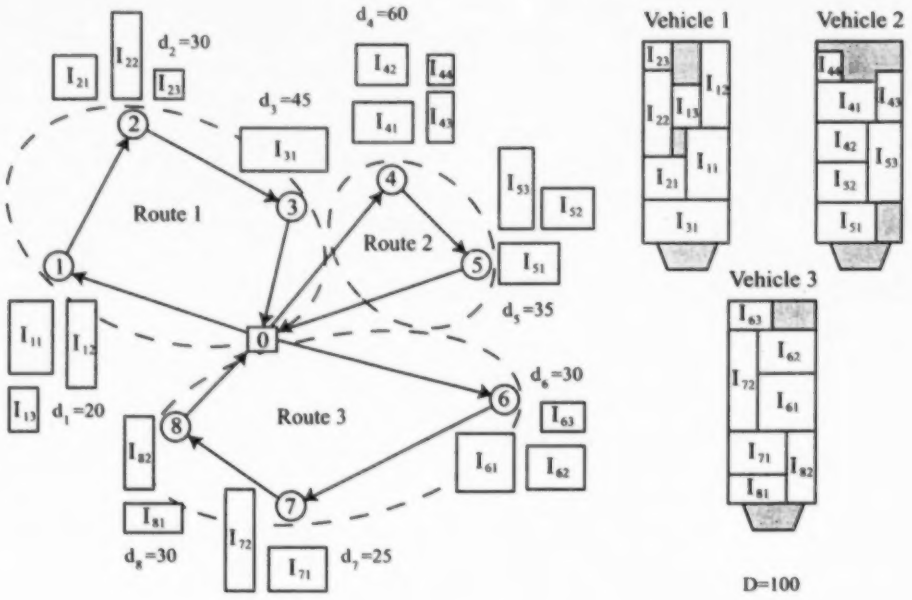

Figure 1: *Example of the 2L-CVRP.*

We will consider two variants of the problem: the *Unrestricted* 2L-CVRP and the *Sequential* 2L-CVRP. In the former variant no additional constraint is imposed, while in the latter we also impose to each vehicle, say k, the following constraint:

Sequence constraint: the vehicle loading must be such that when client $i \in S(k)$ is visited, the items of his lot, $\bigcup_{\ell \in \{1,\ldots,m_i\}} I_{i\ell}$, can be downloaded through a sequence of straight movements (one per item) parallel to the H-edge of the loading area.

In other words, for any item $I_{i\ell}$, no item demanded by a client visited later on in the tour must lay on the strip going from the w-edge of $I_{i\ell}$ to the rear of the vehicle. The dashed strip above item I_{il} in Figure 2 (a) shows the forbidden loading area.

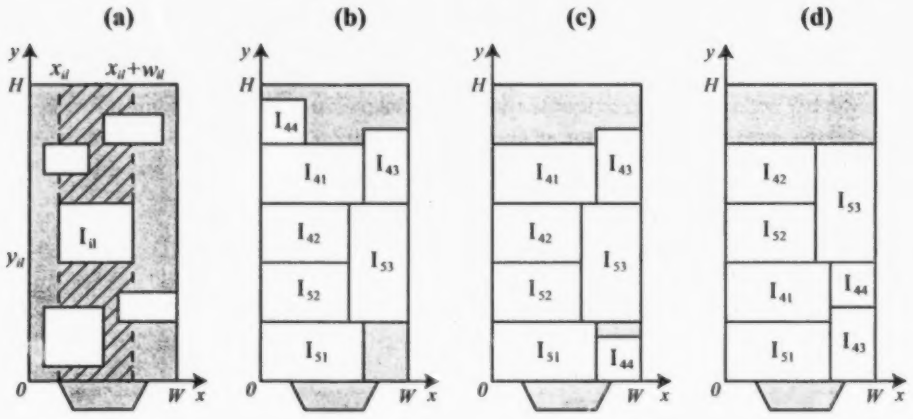

Figure 2: *The loading surface.*
(a) *forbidden areas*; (b) *sequential loading*; (c), (d) *unrestricted loadings.*

This constraint can arise in practice when the items have great weight, size or fragility, so that moving them inside the vehicle is extremely difficult or impossible. The sequential variant is the one studied in [20]. In Figure 1, we show a simple example with 3 vehicles of weight capacity $D = 100$, 8 customers and 21 items. The three patterns given for the vehicles are feasible sequential loadings for the freight delivery. (An example of non-sequential loading is given in Figure 2.)

We assume in the following that all input data are positive integers. The loading area of a vehicle is represented in the positive quadrant of a Cartesian coordinate system, with its bottom left corner in $(0, 0)$ and the W-edge (resp. H-edge) parallel to the x-axis (resp. y-axis), and the vehicle rear coinciding with segment $[(0, H), (W, H)]$. The position of an item $I_{i\ell}$ in a loading pattern is given by the coordinates of its bottom left point, $(x_{i\ell}, y_{i\ell})$. The coordinate system is shown in Figure 2: by assuming that $W = 100$ and $H = 200$, in Figure 2 (b) the position of item I_{51} is $(0, 0)$, that of item I_{52} is $(0, 40)$, and so on.

3 A Tabu Search Heuristic

Tabu Search is a very effective and simple technique for the solution of optimization problems. It was first proposed by Glover [14, 15] and has since been applied to a large number of different problems (see, e.g., Glover [16] and Glover and Laguna [17] for surveys).

Tabu Search starts from an initial (usually feasible) solution and iteratively moves to a new one selected in a certain *neighborhood* of the current solution. At each iteration, the *move* yielding the best solution in the neighborhood is selected, even if this results in a worse solution. To avoid cycling, solutions possessing some attributes of previously visited solutions are declared tabu for a number of iterations. In order to obtain good practical results, this simple scheme is usually enriched with additional features. An *intensification* phase is often performed to accentuate the search in a promising area, while *diversification* is necessary to escape from a poor or already extensively searched area.

The Tabu Search approach introduced in the next sections can accept moves producing infeasible tours in the following sense. For the sequential 2L-CVRP, all tours considered in the search satisfy the sequence constraint. However, both in the sequential and the unrestricted 2L-CVRP, the tours can be *weight-infeasible* if the total weight exceeds D, or *load-infeasible* if the packing needs a loading surface of height exceeding H. The algorithm never produces packings requiring a loading surface wider than W or implying overlapping items. Infeasible moves are assigned a *penalty* proportional to the level of the violation.

Thus, when performing a move, the algorithm must consider the improvement of the current solution in terms of total edge cost, and check the feasibility of the candidate tour. For the routing aspects, we have adopted an approach (described later in Section 3.3) derived from some successful features of Taburoute, a Tabu Search heuristic developed by Gendreau, Hertz and Laporte [12] for the solution of vehicle routing problems. For what concerns feasibility, the weight constraint is immediate to check, but the loading constraint requires the solution of an NP-hard problem. We have thus developed a heuristic algorithm, described in the next section, executed for each set of clients assigned to a vehicle: it outputs a two-dimensional pattern of width W and height $Hstrip$, to be tested against the available height H.

3.1 A heuristic algorithm to check the loading constraints

Let $I(k)$ (see (1)) be the set of items currently assigned to vehicle k. In this subsection we describe two heuristics, LH_{2SL} and LH_{2UL}, to check the loading constraints for vehicle k in the sequential and the unrestricted case, respectively. The algorithms iteratively apply an inner procedure based on the heuristic rule used in algorithm TP_{RF} (*Touching Perimeter*) proposed by Lodi, Martello and Vigo [21] for the two-dimensional bin packing problem, and applied by Iori, Martello and Monaci [19] to the two-dimensional strip packing problem. We have developed two inner procedures, TP_{2SL} and TP_{2UL}, for the sequential and the unrestricted case, respectively. Both procedures assign one item at a time (according to a given item sorting) to a strip of width W, trying to maximize at each iteration the percentage of the item perimeter touching the strip or other items already packed (*touching*

perimeter). Each item is packed in the *normal position* (as defined by Christofides and Whitlock [3]), i.e., with its bottom edge touching either the bottom of the strip or the top edge of another item, and with its left edge touching either the left edge of the strip or the right edge of another item.

We next provide the description of procedure TP_{2SL} (sequential case), shown in Figure 3. The first call to TP_{2SL} is performed by LH_{2SL} with an item sorting obtained by scanning the clients assigned to a route in reverse order of visit and, for each client, by listing the items in the corresponding lot according to non-increasing width, breaking ties by non-increasing height. Subsequent calls are performed by perturbing this order, as will be shown. Procedure TP_{2SL} can terminate in one of three possible ways: with a feasible packing, with a load-infeasible packing on a surface $W \times Hstrip$ with $Hstrip > H$ (to be penalized within Tabu Search), or with no packing at all. Since the last outcome cannot occur at the first call, procedure LH_{2SL} always returns a feasible or load-infeasible packing.

procedure $\mathbf{TP_{2SL}}(I(k))$
1. pack the first item of $I(k)$ in position $(0,0)$ and set $\overline{H}$ at its height;
2. **for each** successive item, say I of size (w,h), of $I(k)$ **do**
   2.1 P:= {set of normal packing positions (x,y) : (a), (b), (c) below hold}
   (a) $x + w \leq W$;
   (b) the rectangle of sides $([(x,y),(x+w,y)], [(x,y),(x,y+h)])$ is free;
   (c) for each already packed item $\overline{I}$, say of size $(\overline{w},\overline{h})$ packed at $(\overline{x},\overline{y})$,
   (c.1) if $\overline{I}$ belongs to the lot of a previous client then
   either $(\overline{y} \geq y)$ or $(\overline{x} + \overline{w} \leq x)$ or $(\overline{x} \geq x + w)$;
   (c.2) if $\overline{I}$ belongs to the lot of a successive client then
   either $(\overline{y} \leq y + h)$ or $(\overline{x} + \overline{w} \leq x)$ or $(\overline{x} \geq x + w)$;
   2.2 **if** $P = \emptyset$ **then return** ∞ **else** $\overline{P} = \{(x,y) \in P : y + h \leq H\}$;
   2.3 **if** $\overline{P} \neq \emptyset$ **then**
   pack I in the position $(x,y) \in \overline{P}$ producing maximum touching perimeter
   else pack I in the position $(x,y) \in P$ for which $y + h$ is a minimum;
   2.4 $\overline{H} := \max\{\overline{H}, y + h\}$
   end for;
3. **return** $\overline{H}$.

Figure 3: *Inner heuristic for producing sequential loadings.*

The procedure works as follows. The first item is packed in position $(0,0)$. The next item, say $I_{i\ell}$, is packed, if possible, with its bottom-left corner at a coordinate (x,y) that maximizes the touching perimeter, provided that (i) the unloading sequence is feasible (see Figure 2 (a)), (ii) the induced strip height does not exceed the vehicle limit. If no such coordinate exists, the item is packed in that position, among those satisfying (i) (if any), for which the induced strip height is smaller. With item sortings other than the initial

one it is possible that an item is encountered for which no packing satisfying (i) exists, in which case a dummy infinite height is returned. Figure 2 (b) shows the packing produced by TP_{2SL} for the second vehicle of Figure 1 with the initial sorting.

Let $\overline{m}$ be the total number of items required by the clients of the considered route. For each item, the number of candidate normal packing positions (set P) is $O(\overline{m})$. The feasibility check of Step 2.1 (c) and the touching perimeter computation of step 2.3 require $O(\overline{m})$ time. The overall time complexity of TP_{2SL} is thus $O(\overline{m}^3)$.

Procedure TP_{2SL} is iteratively invoked by algorithm LH_{2SL} on modified item sequences as shown in Figure 4, where λ is the maximum number of times the inner procedure is executed. On input, the items are ordered according to the initial sorting previously described. The output is either a feasible packing if the returned value is $Hstrip \leq H$, or a load-infeasible packing if $Hstrip > H$. Since Procedure TP_{2SL} finds one of these two kinds of packing when invoked with the initial sorting, LH_{2SL} will always return a packing.

Procedure $LH_{2SL}(I(k),\lambda)$
1. $Hstrip := TP_{2SL}(I(k))$, $it := 1$;
2. **while** $Hstrip > H$ **and** $it \leq \lambda$ **do**
   2.1 randomly select two items and switch their positions;
   2.2 $\xi := TP_{2SL}(I(k))$;
   2.3 $Hstrip := \min\{Hstrip, \xi\}$;
   2.4 $it := it + 1$;
   end while
3. **return** $Hstrip$.

Figure 4: *Heuristic algorithm for checking the sequential loading constraints.*

For the unrestricted case, the corresponding procedure, LH_{2UL}, is identical to LH_{2SL} except for the call to a procedure TP_{2UL} which is obtained from TP_{2SL} by dropping condition (c) of Step 2.1. (Note that in this case P will never be empty at Step 2.2.) For the vehicle of Figure 2 (b), Figures 2 (c) and (d) show two packings produced by TP_{2UL}: the former with the initial sorting, the latter with items sorted by decreasing width, breaking ties by decreasing height.

3.2 Initial solution

We use two heuristic algorithms (each in two versions, for the sequential and the unrestricted cases) for determining an initial solution. The first algorithm, IHG_{2SL} or IHG_{2UL}, works for general 2L-CVRP instances and is always executed. For the special case of Euclidean instances, a second heuristic algorithm, IHE_{2SL} or IHE_{2UL}, is also executed.

Algorithm IHG_{2SL} (resp. IHG_{2UL}) was obtained by embedding the loading and sequencing constraints (resp. the loading constraints) in the classical Clarke and Wright [4]

savings algorithm. Whenever an attempt is made to merge two routes, we check whether the weight constraint is satisfied, and, if so, we execute procedure LH_{2SL} (resp. LH_{2UL}) of Section 3.1 to check the load feasibility. The algorithm is first executed by only accepting merges for which a feasible packing is returned. If in this way we obtain a solution that uses v vehicles, the execution terminates. Otherwise, starting from the current solution, we execute a new round where load-infeasible packings are accepted.

The Euclidean algorithm, IHE_{2SL} (resp. IHE_{2UL}), was obtained by embedding the loading and sequencing constraints (resp. the loading constraints) in the algorithm developed by Cordeau, Gendreau and Laporte [6] for periodic and multi-depot VRPs. We start by randomly selecting a radius emanating from the depot, and by sorting the clients by increasing value of the angle between the radius and the segment connecting the client to the depot. The first client is assigned to the first vehicle. We then attempt to assign each subsequent client to the current vehicle by checking weight capacity and loading feasiblity through LH_{2SL} (resp. LH_{2UL}). For the first $v-1$ vehicles, only feasible packings are accepted, while a load-infeasible packing is only accepted for the last vehicle. The best feasible solution determined in the initialization phase (if any) is stored as the incumbent solution.

3.3 Tabu Search

As previously observed, our Tabu Search heuristic has to deal with weight-infeasible or load-infeasible solutions. We treat infeasibilities as penalties in an objective function to be minimized. Let s denote a solution, consisting of a set of $\tilde{v}$ routes, with $1 \leq \tilde{v} \leq v$, in which each client belongs to exactly one route. Let $c(k)$ be the total cost of the edges in route k. As in Taburoute, we express the objective function as a sum of three terms:

$$z'(s) = z(s) + \alpha q(s) + \beta h(s), \tag{2}$$

where $q(s)$ and $h(s)$ (see below) represent the entity of the violation, if any, of the weight and loading constraints, respectively, while α and β are self-adjusting parameters. Using the notation introduced in Section 2 and denoting by H_k the height of the two-dimensional loading of vehicle k, we have

$$z(s) = \sum_{k=1}^{\tilde{v}} c(k), \tag{3}$$

$$q(s) = \sum_{k=1}^{\tilde{v}} \left[\sum_{i \in S(k)} d_i - D \right]^+, \tag{4}$$

$$h(s) = \sum_{k=1}^{\tilde{v}} [H_k - H]^+, \tag{5}$$

where $[a]^+ = \max\{0, a\}$.

A move consists of selecting a client i, currently inserted in a route k, removing it from the route and inserting it into a new route k'. The cost of the new solution is obtained through the GENI (Generalized Insertion Procedure) heuristic developed by Gendreau, Hertz and Laporte [13] for the Travelling Salesman Problem. GENI considers a subset of the possible 4-opt modifications of the two tours (the one from which the client is deleted and the one where it is inserted), selecting the best one for each tour. Let $N_p(i)$ be the set of the p vertices closest to i (*neighbors* of client i). In order to reduce the number of possible moves, only empty routes and routes containing at least one neighbor of i are considered for possible insertion of i. A good compromise between computing time and solution quality was experimentally obtained with $p = 10$.

At each iteration, all possible moves are evaluated, i.e., for each client we consider its possible insertion into all routes satisfying the neighbor condition above: for each move, the resulting cost $z'(s)$ (see (2)–(5)) is obtained by determining $z(s)$ through GENI, directly computing $q(s)$, and computing $h(s)$ through procedure $\mathrm{LH}_{\mathrm{2SL}}$ (or $\mathrm{LH}_{\mathrm{2UL}}$). In certain situations, a *penalty term* π (see below, equation (6)) is possibly added to the cost in order to diversify the search. The move producing the solution with minimum overall cost is selected. After a move has been performed by removing client i from route k, reinserting i in k is declared tabu for the next ϑ iterations. An aspiration criterion is based on an $n \times v$ array M that stores in M_{ij} the value of the best feasible solution so far obtained, if any, with client i assigned to route j ($M_{ij} = \infty$ if no such solution has been identified). A tabu move that reinserts client i in route k can be accepted if it improves the current M_{ij} value. The value of ϑ was experimentally determined as $\vartheta = \log(nv)$.

Three main parameters govern the search process: α controls weight constraint violations (see (2)), while β and λ control packing constraint violations (see (2) and Section 3.1). In the initialization phase α and β are set to 1, and λ to 2. In the iterative phase the three parameters are recursively updated following the characteristics of the solutions produced by the moves. If the current solution is (resp. is not) weight-infeasible we set $\alpha = \alpha(1 + \delta)$ (resp. $\alpha = \alpha/(1 + \delta)$). In addition, if the current solution is (resp. is not) load-infeasible we set $\beta = \beta(1 + \delta)$ and $\lambda = \min\{\lambda + 1, 4\}$ (resp. $\beta = \beta/(1 + \delta)$ and $\lambda = \max\{\lambda - 1, 1\}$. A good value of δ was experimentally determined as $\delta = 0.01$.

We use a simple diversification technique to escape from poor local minima. Let s_a be the current solution, and s_b the solution produced by a move that removes customer i from route k. As in Taburoute, if $z'(s_b) \geq z'(s_a)$, we add to $z'(s_b)$ a penalty term π, proportional to the frequency of the considered move during the previous iterations. Let ρ_{ik} denote the number of times customer i was removed from route k, divided by the total number of iterations. The penalty term is then

$$\pi = \rho_{ik}(z'(s_b) - z'(s_a))\gamma, \tag{6}$$

with γ set to $\sqrt{nv}$, as suggested by Taillard [26] and Cordeau, Laporte and Mercier [8]. In this way, when no downhill move is found, the algorithm is more likely to escape from poor local search areas. In order to deeply explore promising areas of the solution space, we extensively use intensification algorithms. Because of their importance they are described in the next section.

4 Intensification

Two kinds of intensification techniques are adopted, one executed periodically, every ω moves, and one (computationally heavier) executed whenever a weigth-feasible potential new incumbent is obtained. Both intensification processes make use of a procedure ($\text{Repack}_{2\text{SL}}(k)$ for the sequential case, or $\text{Repack}_{2\text{UL}}(k)$ for the unrestricted case) that looks for a feasible packing for the items of each route k of a load-infeasible solution s, i.e., one for which $H_k > H$.

Recall that $I(k)$ is the set of items associated with the clients of route k. Procedure $\text{Repack}_{2\text{SL}}(I(k), \lambda)$, shown in Figure 5, starts by computing lower bounds (Martello and Vigo [24]) for the two-dimensional bin packing problem instance induced by the items associated with route k and bin size $W \times H$. If more than one bin is needed, no feasible loading exists and the procedure terminates. Otherwise we execute the heuristic search $\text{LH}_{2\text{SL}}$ of Section 3.1. For this computation, it turned out to be convenient to allow a much higher number of calls, $\lambda = 10$, to the inner heuristic $\text{TP}_{2\text{SL}}$. If no feasible packing is still found, we execute the branch-and-bound algorithm of Iori, Salazar González and Vigo [20] with a limit of 1 000 backtrackings.

Procedure $\text{Repack}_{2\text{SL}}(I(k), \lambda)$
1. $L :=$ lower bound for the 2BPP instance defined by $I(k)$ with bin size $W \times H$;
2. **if** $L > 1$ **then return** ∞;
3. $\xi := \text{LH}_{2\text{SL}}(I(k), \lambda)$;
4. **if** $\xi \leq H$ **then return** ξ;
5. execute a truncated branch-and-bound search;
6. **if** a solution of value 1 is found **then return** H
   else return ∞.

Figure 5: *Intensification for packing constraints.*

Procedure $\text{Repack}_{2\text{UL}}$ is identical to $\text{Repack}_{2\text{SL}}$ except for the call to $\text{LH}_{2\text{UL}}$ instead of $\text{LH}_{2\text{SL}}$, and for the execution of a truncated branch-and-bound search that does not consider the sequence constraints.

4.1 Periodic Intensification

Every ω moves, a local search for improving the incumbent solution is performed through a modified version of the heuristic algorithm of Iori, Salazar González and Vigo [20]. This algorithm is embedded within a branch-and-cut approach and operates on a modified cost matrix $\bar{c}$ obtained by setting, for each edge (i, j), $\bar{c}_{ij} = c_{ij}(1 - x_{ij}/2)\mu$, where x_{ij} is the (possibly fractional) value of the 0-1 decision variable associated with the edge and μ is a randomization factor uniformly distributed in [0,1]. The new solution is then obtained by iteratively building routes through a parametric version of the savings algorithm. Each

route (client set) S is constructed by first selecting the client i for which the quantity $\sigma_c(c_{0i}+c_{i0})+\sigma_d d_i+\sigma_a a_i$ is a maximum (σ_c, σ_d and σ_a prefixed parameters), and then iteratively inserting in S the client j for which the residual weight capacity $d_{res}=D-\sum_{i\in S} d_i-d_j$ and loading area $a_{res}=A-\sum_{i\in S} a_i-a_j$ are non-negative, and σ_c(insertion cost) $+\sigma_d d_{res}+\sigma_a a_{res}$ is a minimum. Our implementation, procedure Improve$_{2SL}$, is shown in Figure 6.

Procedure Improve$_{2SL}$
1. create the modified cost matrix $\bar{c}$ and set $P:=\emptyset$;
2. **for each** triplet $\sigma=(\sigma_c,\sigma_d,\sigma_a)\in T$ **do**
   $s:=$ SAVINGS$(\bar{c},\sigma)$;
   if s includes no more than v routes **then**
   if $z(s)<z$ **then** $P:=P\cup\{s\}$;
   $z_{\text{old}}:=\infty$;
   while $z(s)<z_{\text{old}}$ **do**
   $z_{\text{old}}:=z(s)$;
   $s:=$ TWO-OPT(c,s);
   if $z(s)<z$ **then** $P:=P\cup\{s\}$
   end while
   end if
   end for.
3. sort the solutions $s\in P$ by non-decreasing $z(s)$ value;
4. **for each** $s\in P$ in order **do**
   for each route k of s **do**
   $\xi:=$ Repack$_{2SL}(I(k),\lambda)$;
   if $\xi>H$ **then break**
   end for;
   if $\xi\le H$ **then return** s
   end for;
   return $\emptyset$.

Figure 6: *Algorithm for improving the incumbent solution.*

The modified cost matrix is obtained by setting

$$\bar{c}_{ij}=\begin{cases} rc_{ij}\ (\text{with } r \text{ uniformly random in } [0,2]) & \text{if } (i,j) \text{ belongs to a route,} \\ c_{ij} & \text{otherwise.} \end{cases} \tag{7}$$

and only clients for which the residual loading area a_{res} is no less than a threshold value τ are considered for possible insertion, in order to take into account the difficulty of packing the corresponding items. The value τ is randomly generated in $[0, A\cdot R/4]$, where

$$R=\max\{\max_{i\ell}\{h_{i\ell}/H\},\max_{i\ell}\{w_{i\ell}/W\}\} \tag{8}$$

(maximum relative size of an item) measures the "difficulty" of the packing problem associated with the instance. In other words, we decrease the residual packing area by a percentage that takes into account the probability that portions of the loading surface remain unused.

We use a set of five triplets for the values of $(\sigma_c, \sigma_d, \sigma_a)$: $T = \{(1,0,0), (\frac{1}{2}, \frac{1}{2}, 0), (\frac{1}{2}, 0, \frac{1}{2}), (0, \frac{1}{2}, \frac{1}{2}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})\}$. The savings algorithm is executed for each triplet, and all solutions that use no more than v routes are improved through 2-opt modifications. In addition, all solutions of value less than that of the incumbent are stored in a *pool* P. We then consider the solutions in P in increasing value. For each solution, we try to obtain a feasible packing for each route through procedure $\text{Repack}_{2SL}(I(k), \lambda)$, stopping as soon as an overall feasible solution (if any) is obtained. Given $\bar{c}$ and a triplet σ, let $s = \text{SAVINGS}(\bar{c}, \sigma)$ be the solution built by the savings algorithm, and TWO-OPT(c,s) the improved solution produced by the first improvement (if any) found through the two-opt local search algorithm (TWO-OPT$(c, s) \equiv s$ if no improvement was found). All solution values $z(s)$ are obviously evaluated using the original cost matrix c.

Procedure Improve_{2UL} is identical to Improve_{2SL} except for the call to $\text{Repack}_{2UL}(I(k), \lambda)$ instead of $\text{Repack}_{2SL}(I(k), \lambda)$. Basing on the outcome of computational testings, we set ω to 1 for $n \leq 40$ and to 5 for larger values.

4.2 Special intensification

Whenever a solution s having a cost $z(s)$ lower than that of the incumbent is produced, different actions are taken, depending on the feasibility of the weight and of the loading associated with s:

(i) if s is weight-infeasible, no attempt is made to make it feasible;

(ii) if s is totally feasible, the incumbent is updated, and some more time is spent in an attempt to further improve s on an enlarged neighborhood. This is achieved by doubling, for each client i, the size of $N_p(i)$, and applying GENI for each resulting possible move. In our implementation this results in considering all routes that are empty or contain at least one of the 20 vertices closest to i;

(iii) if s is weight-feasible but load-infeasible, we first iteratively execute procedure Repack for each route k of s for which $H_k > H$. If a feasible solution is obtained in this way, the additional intensification of case (ii) above is further performed.

5 Computational Results

The Tabu Search heuristic developed in the previous sections was coded in C and tested on a test set obtained by modifying instances from the literature. All experiments were executed on a Pentium IV 1700 MHz. The graph and the weights demanded by the customers were obtained from CVRP instances (test instances for the CVRP are described in [27], and

can be downloaded from `http://www.or.deis.unibo.it/research.html`). The number of items and their dimensions were created according to the five classes introduced in [20] and [18], namely:

Class 1: each customer i, for $i = 1, \ldots, n$ is assigned one item with unit width and height.

Classes 2 – 5: a random uniform distribution on a certain interval (see Table 1, column 2) is used to generate the number m_i of items for each customer i. Each item is then randomly assigned, with equal probability, one out of three possible shapes (*vertical*, *homogeneous* or *horizontal*). Finally, the dimensions of the items are uniformly randomly generated in a given interval (see again Table 1, columns 3–8).

The values $H = 40$ and $W = 20$ were chosen for the dimensions of the loading area. Class 1 corresponds to pure CVRP instances, for which no loading constraint is imposed. The other classes proved to be a challenging mix of items of different sizes.

Table 1: Classes 2 –5

		Vertical		*Homogeneous*		*Horizontal*	
Class	m_i	h_{il}	w_{il}	h_{il}	w_{il}	h_{il}	w_{il}
2	$[1, 2]$	$\left[\frac{4H}{10}, \frac{9H}{10}\right]$	$\left[\frac{W}{10}, \frac{2W}{10}\right]$	$\left[\frac{2H}{10}, \frac{5H}{10}\right]$	$\left[\frac{2W}{10}, \frac{5W}{10}\right]$	$\left[\frac{H}{10}, \frac{2H}{10}\right]$	$\left[\frac{4W}{10}, \frac{9W}{10}\right]$
3	$[1, 3]$	$\left[\frac{3H}{10}, \frac{8H}{10}\right]$	$\left[\frac{W}{10}, \frac{2W}{10}\right]$	$\left[\frac{2H}{10}, \frac{4H}{10}\right]$	$\left[\frac{2W}{10}, \frac{4W}{10}\right]$	$\left[\frac{H}{10}, \frac{2H}{10}\right]$	$\left[\frac{3W}{10}, \frac{8W}{10}\right]$
4	$[1, 4]$	$\left[\frac{2H}{10}, \frac{7H}{10}\right]$	$\left[\frac{W}{10}, \frac{2W}{10}\right]$	$\left[\frac{H}{10}, \frac{4H}{10}\right]$	$\left[\frac{W}{10}, \frac{4W}{10}\right]$	$\left[\frac{H}{10}, \frac{2H}{10}\right]$	$\left[\frac{2W}{10}, \frac{7W}{10}\right]$
5	$[1, 5]$	$\left[\frac{H}{10}, \frac{6H}{10}\right]$	$\left[\frac{W}{10}, \frac{2W}{10}\right]$	$\left[\frac{H}{10}, \frac{3H}{10}\right]$	$\left[\frac{W}{10}, \frac{3W}{10}\right]$	$\left[\frac{H}{10}, \frac{2H}{10}\right]$	$\left[\frac{W}{10}, \frac{6W}{10}\right]$

Some details on the instances generated in this way are reported in Table 2. For each CVRP instance, we created five 2L-CVRP instances (one per class) by generating the number of vehicles v through heuristic $\mathrm{LH}_{\mathrm{2SL}}$ of Section 3.1, modified so as to pack the items into finite $W \times H$ bins instead of a strip. The CVRP instances 1–16 were already used in [20] and [18], producing 80 different 2L-CVRP instances according to the five above classes, while the 100 instances obtained from 17–36 are considered here for the first time. In the lines of Table 2 (one for each group of five instances), the first column gives the name of the original CVRP instance and the second one the number of customers (n). The five following groups of three columns give, for each class, the total number of items ($M = \sum_{i=1}^{n} m_i$), the number of vehicles (v) and a lower bound (LB) on the minimum number of vehicles needed to ensure feasibility. This lower bound, reported here in order to show that the produced instances are indeed "reasonable", was computed as the maximum among two valid lower bounds: the number of vehicles in the original CVRP instance and the value obtained by solving the two-dimensional bin packing problem induced by the item

Table 2: Details on the instances. Original CVRP instance, number of customers; for each class, number of items, number of vehicles, lower bound on a feasible number of vehicles.

		Class 1			*Class 2*			*Class 3*			*Class 4*			*Class 5*		
Instance	n	M	v	LB	M	v	LB	M	v	LB	M	v	LB	M	v	LB
1) E016-03m	15	15	3	3	24	3	3	31	3	3	37	4	3	45	4	3
2) E016-05m	15	15	5	5	25	5	5	31	5	5	40	5	5	48	5	5
3) E021-04m	20	20	4	4	29	5	4	46	5	4	44	5	4	49	5	4
4) E021-06m	20	20	6	6	32	6	6	43	6	6	50	6	6	62	6	6
5) E022-04g	21	21	4	4	31	4	4	37	4	4	41	4	4	57	5	4
6) E022-06m	21	21	6	6	33	6	6	40	6	6	57	6	6	56	6	6
7) E023-03g	22	22	3	3	32	5	4	41	5	4	51	5	4	55	6	3
8) E023-05s	22	22	5	5	29	5	5	42	5	5	48	5	5	52	6	5
9) E026-08m	25	25	8	8	40	8	8	61	8	8	63	8	8	91	8	8
10) E030-03g	29	29	3	3	43	6	5	49	6	4	72	7	6	86	7	5
11) E030-04s	29	29	4	4	43	6	5	62	7	6	74	7	6	87	7	5
12) E031-09h	30	30	9	9	50	9	9	56	9	9	82	9	9	101	9	9
13) E033-03n	32	32	3	3	44	7	5	56	7	5	78	7	6	102	8	5
14) E033-04g	32	32	4	4	44	7	5	57	7	5	65	7	5	87	8	4
15) E033-05s	32	32	5	5	50	6	5	59	6	6	90	8	7	114	8	6
16) E036-11h	35	35	11	11	56	11	11	74	11	11	93	11	11	114	11	11
17) E041-14h	40	40	14	14	60	14	14	73	14	14	96	14	14	127	14	14
18) E045-04f	44	44	4	4	66	9	7	87	10	8	112	10	8	122	10	6
19) E051-05e	50	50	5	5	82	11	9	103	11	10	134	12	10	157	12	8
20) E072-04f	71	71	4	4	104	14	12	151	15	13	178	16	13	226	16	13
21) E076-07s	75	75	7	7	114	14	12	164	17	14	168	17	14	202	17	14
22) E076-08s	75	75	8	8	112	15	13	154	16	14	198	17	14	236	17	14
23) E076-10e	75	75	10	10	112	14	13	155	16	14	179	16	14	225	16	14
24) E076-14s	75	75	14	14	124	17	14	152	17	14	195	17	14	215	17	14
25) E101-08e	100	100	8	8	157	21	18	212	21	18	254	22	19	311	22	19
26) E101-10c	100	100	10	10	147	19	16	198	20	17	247	20	18	310	20	18
27) E101-14s	100	100	14	14	157	19	17	211	22	19	245	22	19	320	22	19
28) E121-07c	120	120	7	7	183	23	20	242	25	21	299	25	21	384	25	21
29) E135-07f	134	134	7	7	197	24	21	262	26	22	342	28	24	422	28	24
30) E151-12b	150	150	12	12	225	29	25	298	30	27	366	30	27	433	30	27
31) E200-16b	199	199	16	16	307	38	33	402	40	35	513	42	37	602	42	37
32) E200-17b	199	199	17	17	299	38	33	404	39	34	497	39	34	589	39	34
33) E200-17c	199	199	17	17	301	37	32	407	41	35	499	41	36	577	41	36
34) E241-22k	240	240	22	22	370	46	40	490	49	42	604	50	44	720	50	44
35) E253-27k	252	252	27	27	367	45	39	507	50	43	634	50	45	762	50	45
36) E256-14k	255	255	14	14	387	47	41	511	51	44	606	51	44	786	51	44

sizes and the vehicle loading area. This second part of the bound was computed through the code by Martello, Pisinger and Vigo [23], available at `http://www.diku.dk/~pisinger`. (This code works for the three-dimensional bin packing problem and can obviously be used for the two-dimensional case as well.) The table shows that the v and LB values are very close: out of a total of 180 instances, these values are identical in 73 cases, and differ by at most two units in 57 cases. The complete set of 180 test instances that were obtained may be downloaded from `http://www.or.deis.unibo.it/research.html`.

In Tables 3 and 4, the performance of our approach is compared with the branch-and-cut algorithm by Iori, Salazar González and Vigo [20] on small-size instances involving less than 45 customers. Since the latter approach does not consider single-customer routes and always uses the v available vehicles, the Tabu Search algorithm was executed by making infeasible (through penalties) any solution not satisfying these properties. In addition, as in [20], the edge costs were computed by rounding down to the next integer the Euclidean distances between vertex pairs. The branch-and-cut algorithm had a time limit of 24 hours per instance, as in the experiments presented in [20]. A convenient policy for controlling the Tabu Search was experimentally determined as follows. The algorithm is halted after $2n^2v$ iterations, or after one hour of CPU time.

For each instance, Tables 3 and 4 give the best solution obtained both by branch-and-cut (z_{bc}) and Tabu Search (z_{TS}), the elapsed CPU time in seconds when the best solution was found ($seconds_h$), and the total CPU time in seconds ($seconds$). Because the branch-and-cut is not always run to completion, the value of z_{bc} may be higher than that of z_{TS}. In addition, asterisks indicate proven optimal solutions, and $\%gap$ evaluates the quality of the Tabu Search solution as $\%gap = 100\ (z_{TS} - z_{bc})/z_{bc}$. The solutions found by the Tabu Search algorithm are very close to optimality, and are provably optimal in 33 out of 57 known optima. For half of the instances with 29 customers or more (Table 4), the Tabu Search solution is better than that found by the branch-and-cut in a much higher CPU time (negative $\%gap$ values). On the other hand, for about one third of the instances with 25 customers or less (Table 3) the Tabu Search CPU time is higher than that of branch-and-cut and, in six of these cases, no optimal solution is found. The worst percentage error is 7.69, the best percentage improvement is 21.23. On average the Tabu Search solutions improve those obtained through branch-and-cut by 0.21%. The CPU times required by the Tabu Search algorithm are very reasonable. On average it takes 116 seconds in total, and 43 seconds to find the best solution, while for the branch-and-cut algorithm these values increase to 29 858 and 12 875 seconds, respectively.

In Table 5 we summarize the results for the complete set of instances, both for sequential and unrestricted loading. The limit given to the Tabu Search heuristic is again the first occurrence between $2n^2v$ iterations and one hour of CPU time. The edge costs are computed as the rational Euclidean distances between any pair of vertices, without rounding them to integer values. One-customer routes are allowed, and fewer than v vehicles may be used. The entries in the second column give the solution value obtained for Class 1, i.e., for pure CVRP instances. In the remaining columns we give, respectively for the sequential and unrestricted case, the average solution values for Classes 2–5 ($\overline{z}_{2-5}$), the average elapsed CPU time when the final solution was found ($seconds_h$) and the average total CPU time

Table 3: Performance of the Tabu Search heuristic with respect to branch-and-cut. Sequential instances with $n \leq 25$, integer edge costs.

Instance		*Branch-and-cut*				*Tabu Search*			
no	*cl*	z_{bc}		$seconds_h$	*seconds*	z_{TS}	$seconds_h$	*seconds*	*%gap*
1	1	273	*	3.80	3.86	273	0.05	2.45	0.00
	2	285	*	51.38	68.81	285	0.19	4.56	0.00
	3	280	*	17.06	21.42	280	1.33	7.84	0.00
	4	288	*	1.80	2.36	290	0.30	7.53	0.69
	5	279	*	53.08	53.08	279	2.28	15.75	0.00
2	1	329	*	0.19	0.55	329	0.09	1.36	0.00
	2	342	*	11.08	11.95	342	1.12	1.98	0.00
	3	347	*	6.39	8.14	350	0.17	3.17	0.86
	4	336	*	22.42	23.42	336	0.06	6.70	0.00
	5	329	*	29.22	29.36	329	0.34	6.08	0.00
3	1	351	*	14.92	15.61	351	0.25	8.36	0.00
	2	389	*	32.59	65.78	407	8.95	9.77	4.63
	3	387	*	6.03	6.06	387	0.97	19.95	0.00
	4	374	*	39.28	39.36	374	14.28	21.94	0.00
	5	369	*	0.17	0.22	369	0.95	29.88	0.00
4	1	423	*	0.38	0.41	423	0.16	5.66	0.00
	2	434	*	5.30	5.47	434	0.27	7.41	0.00
	3	432	*	5.23	8.16	438	5.03	12.77	1.39
	4	438	*	23.95	26.38	451	0.81	16.91	2.97
	5	423	*	23.00	44.70	423	1.16	27.13	0.00
5	1	367	*	0.08	0.08	367	2.94	12.83	0.00
	2	380	*	7.98	8.36	396	4.00	19.69	4.21
	3	373	*	1.81	1.91	377	0.58	20.55	1.07
	4	377	*	50.20	50.30	406	5.92	33.00	7.69
	5	389	*	2928.20	2928.27	389	5.73	57.84	0.00
6	1	488	*	5.89	10.69	488	0.13	10.30	0.00
	2	491	*	145.83	145.86	498	1.84	13.98	1.43
	3	496	*	135.41	150.37	496	11.25	16.95	0.00
	4	489	*	14.12	16.56	503	0.14	40.17	2.86
	5	488	*	10.76	13.34	488	0.70	34.19	0.00
7	1	558	*	0.02	0.02	558	0.31	22.02	0.00
	2	724	*	32.05	32.47	752	0.67	18.89	3.87
	3	698	*	3.19	4.39	704	19.86	29.78	0.86
	4	714	*	2596.83	2597.27	742	26.55	50.36	3.92
	5	742	*	738.94	747.17	743	16.06	75.08	0.13
8	1	657	*	0.03	0.05	657	6.99	31.33	0.00
	2	720	*	75.94	91.81	720	3.06	20.09	0.00
	3	730	*	70.03	72.95	752	20.55	33.38	3.01
	4	701	*	7.44	14.12	722	11.22	49.95	3.00
	5	721	*	1128.86	1128.86	736	12.05	90.17	2.08
9	1	609	*	6.17	31.91	609	1.89	11.87	0.00
	2	612	*	453.52	460.52	612	2.89	15.44	0.00
	3	615	*	164.83	194.31	626	7.89	38.45	1.79
	4	626	*	852.14	1593.34	627	5.87	43.47	0.16
	5	609	*	47.39	69.17	609	8.47	81.86	0.00

Table 4: Performance of the Tabu Search heuristic with respect to branch-and-cut. Sequential instances with $25 < n \leq 40$, integer edge costs.

Instance		*Branch-and-cut*				*Tabu Search*			
no	*cl*	z_{bc}		$seconds_h$	*seconds*	z_{TS}	$seconds_h$	*seconds*	*%gap*
10	1	524		1226.17	86400.52	544	19.91	97.33	3.82
	2	774		48373.34	86401.15	703	3.81	72.20	−9.17
	3	638	*	2304.46	3149.99	676	47.86	118.67	5.96
	4	738	*	12671.31	12696.29	773	47.28	156.94	4.74
	5	706		48220.23	70308.02	724	84.47	308.87	2.55
11	1	500	*	0.13	0.14	500	0.81	107.81	0.00
	2	789		99.80	86400.46	734	4.69	72.23	−6.97
	3	763		63747.77	86400.32	785	72.17	101.69	2.88
	4	881		85181.52	86400.83	877	196.44	209.52	−0.45
	5	695		64392.94	86400.06	696	279.17	387.00	0.14
12	1	599		50093.46	86400.34	598	19.41	33.80	−0.17
	2	625		75.60	86400.67	628	9.64	42.87	0.48
	3	597		36171.03	86400.58	597	12.91	50.75	0.00
	4	624		86321.63	86400.16	640	95.98	120.59	2.56
	5	602		22352.36	86400.36	597	103.47	188.24	−0.83
13	1	1991	*	13.18	14.52	1991	47.36	218.69	0.00
	2	3523		10784.11	86400.30	2775	43.75	123.06	−21.23
	3	2570	*	3087.87	32124.54	2696	121.50	170.34	4.90
	4	2673		34999.17	86400.38	2743	29.31	277.05	2.62
	5	2807		83628.40	86400.10	2737	237.52	691.88	−2.49
14	1	827		25734.79	86400.43	823	21.76	145.03	−0.48
	2	1459		80521.15	86400.84	1266	52.33	144.38	−13.23
	3	1211		42204.75	86400.09	1204	137.94	207.09	−0.58
	4	1166	*	25391.47	25542.61	1187	108.52	324.94	1.80
	5	1504		80511.08	86400.08	1309	103.17	895.02	−12.97
15	1	907	*	17.81	17.84	907	51.92	196.37	0.00
	2	1203		386.89	86401.44	1135	94.66	133.89	−5.65
	3	1405		15880.73	86401.21	1183	37.11	205.87	−15.80
	4	1358		55614.54	86400.27	1372	268.25	332.20	1.03
	5	1390		59867.80	86400.10	1361	651.45	671.70	−2.09
16	1	682	*	7086.48	7086.71	682	56.06	96.58	0.00
	2	682	*	1766.95	3374.76	682	20.55	91.70	0.00
	3	682	*	345.15	3853.35	682	15.33	138.09	0.00
	4	691	*	33375.93	33391.84	704	21.69	206.08	1.88
	5	682	*	339.03	2783.99	682	62.77	276.58	0.00
17	1	859		493.15	86400.27	842	84.39	158.62	−1.98
	2	866		283.79	86400.41	851	56.61	132.50	−1.73
	3	850		655.98	86400.02	842	38.84	200.61	−0.94
	4	853		102.91	86400.38	845	7.38	279.73	−0.94
	5	845		263.56	86400.52	842	44.73	402.22	−0.36

Table 5: Performance of the Tabu Search heuristic. Real edge costs.

		Unrestricted			*Sequential*		
Inst.	z_1 (CVRP)	$\overline{z}_{2-5}$	$seconds_h$	*seconds*	$\overline{z}_{2-5}$	$seconds_h$	*seconds*
1	278.73	291.60	4.18	9.70	299.09	2.57	9.23
2	334.96	341.02	0.06	3.69	345.23	0.42	3.50
3	359.77	377.35	1.58	19.57	385.30	3.82	18.91
4	430.88	437.45	0.55	14.74	443.42	1.36	16.96
5	375.28	380.20	5.05	25.18	384.06	4.12	27.60
6	495.85	501.02	7.19	18.78	502.78	5.11	19.47
7	568.56	700.34	6.26	48.45	721.90	15.49	53.03
8	568.56	694.99	11.18	61.34	722.73	32.76	83.70
9	607.65	619.69	3.61	41.23	624.06	10.64	39.98
10	538.79	700.39	36.05	191.97	714.90	43.49	179.61
11	505.01	739.04	55.68	207.59	773.45	99.02	199.38
12	610.57	620.62	49.03	82.66	631.85	58.76	99.52
13	2006.34	2598.20	57.46	332.05	2687.03	48.97	312.81
14	837.67	1047.72	375.85	565.63	1101.49	146.03	439.46
15	837.67	1201.38	156.74	375.86	1240.89	165.41	313.45
16	698.61	702.03	20.49	160.15	704.85	28.02	157.17
17	862.62	866.37	64.93	218.21	866.50	88.93	226.21
18	723.54	1085.84	589.30	1317.98	1116.17	566.51	1167.75
19	524.61	772.25	633.66	1524.50	802.48	365.21	1521.47
20	241.97	564.67	954.52	3237.30	581.81	808.86	3370.28
21	688.18	1066.21	460.09	3600.04	1110.19	1702.17	3561.18
22	740.66	1087.46	1191.17	3544.54	1130.33	1573.85	3461.82
23	860.47	1104.72	2032.41	3538.63	1186.36	675.79	3600.03
24	1048.91	1187.62	1454.15	3256.07	1248.43	2642.48	3324.59
25	830.26	1436.09	1205.76	3600.03	1480.63	2336.48	3600.09
26	819.56	1404.49	1173.89	3600.05	1471.74	1554.62	3600.34
27	1099.95	1450.18	521.29	3600.04	1524.22	1308.24	3600.04
28	1078.27	2738.31	2051.16	3600.15	2858.53	2576.93	3600.11
29	1179.01	2474.33	1406.49	3600.05	2575.28	1162.54	3600.19
30	1061.55	1948.72	1185.41	3600.31	2076.20	2021.38	3600.20
31	1464.04	2506.99	2375.77	3600.23	2592.17	2102.18	3600.51
32	1352.61	2486.43	1664.77	3600.44	2605.10	2305.23	3600.55
33	1361.51	2504.00	1843.23	3600.22	2610.55	2221.20	3600.61
34	858.94	1466.06	1359.09	3601.34	1546.06	2184.42	3601.01
35	992.86	1765.30	2061.74	3600.42	1985.44	2223.10	3600.20
36	678.87	1909.88	2265.80	3600.10	1946.66	2626.28	3600.88

(*seconds*). Both times are computed over the five considered classes. (Detailed results for each class can be found in [18]).

The Tabu Search algorithm could find a feasible solution for all 360 instances. By comparing columns z_1 and $\overline{z}_{2-5}$ (unrestricted) one can observe that the inclusion of the loading constraint considerably worsens the solution values of the CVRP (on average by 53.48%). Column $\overline{z}_{2-5}$ (sequential) shows that a relatively smaller further increase (on average by 4.15%) is produced by the constraint on sequential loading. The total CPU times are very close for the two versions, partially also due to the fact that for the largest instances with $n > 75$ the time limit of one hour was almost systematically reached. The time needed to find the best solution is lower for the unrestricted version (average $time_h = 758$ seconds) than for the sequential one (average $time_h = 936$ seconds).

5.1 Robustness

The robustness of the algorithm was tested by running it with different values of the parameters. These tests were performed on the instances of Tables 3 and 4.

The algorithm is quite sensitive to the number of iterations allowed. As previously mentioned, the maximum number of iterations was experimentally determined as $2n^2v$, halting the execution, in any case, after one hour of CPU time. Other rules were tested, leading to a worse average ratio of solution values over computing times. In particular, replacing $2n^2v$ with a fixed limit of value 200 000 improved the average gap from -0.21% to -0.63%, and the number of negative gaps from 19 to 22, but the average computing time increased from 116 to 1547 seconds. Trying 100 000 iterations and 30 minutes halved the CPU time but decreased by one half the average improvement . Attempts with 300 000 iterations did not lead to any significant variation.

The algorithm proved to be particularly robust with respect to the policy adopted for the penalization factor γ, which was finally set to $\sqrt{nv}$, as in [26] and [8]. Other functions attempted, such as $\sqrt{nM}$, nv and nM, produced marginal variations. The approach proved to be also quite robust with respect to the length of the Tabu list, ϑ. The value $\vartheta = \log(nv)$ was finally chosen, but other functions of n and v produced similar results. Functions of the total number of items M proved instead to be a bad choice. As for the size p of the neighborhood, we tested the values 5, 9, 10, 11 and 20. The best choice proved to be $p = 10$. The results were close for $p = 9$ and $p = 11$, but definitely worse for $p = 5$ and $p = 20$.

Different rules were tried for the updating of parameters α and β, but none managed to find the results obtained through the rule described in Section 3.3. The behavior of the algorithm also consistently changed when changing the maximum limit given to λ, the maximum number of calls to $\mathrm{TP}_{\mathrm{2SL}}$ and resp. $\mathrm{TP}_{\mathrm{2UL}}$. The values 1, 2, 3, 4, 5 and 10 were attempted, and the value 4 was chosen. With the value 5, the CPU time increased by 3%, improving the average gap only by 0.02%. An almost symmetric effect was produced by using the value 3.

The periodic intensification (see Section 4.1) is performed every ω iterations. Several values were tested for ω: 10, 5, 4, 3, 2 and 1. It turned out that frequent executions are

useful for small-size instances. The best results were obtained by setting ω to 1 for $n \leq 40$, and to 5 for larger values.

6 Conclusions

We have considered an extension of the classical vehicle routing problem in which two-dimensional packing constraints are introduced. This problem features of two classical combinatorial optimization problems. We have developed, implemented and tested a Tabu Search algorithm which was applied to a large test set obtained by modifying instances from the literature. On instances solved optimally by branch-and-cut, the proposed algorithm also finds an optimal solution in 33 cases out 57. On instances for which the branch-and cut was interrupted before completion, our heuristic finds feasible solutions that are on the average 0.21% better. The Tabu Search algorithm was also applied to a larger set of 360 instances for which the optimum was not known. It succeeded in identifying a feasible solution in all cases. Our results also show that the introduction of a loading constraint considerably increases solution costs.

Acknowledgements

We thank the Ministero dell'Istruzione, dell'Università e della Ricerca (MIUR), the Consiglio Nazionale delle Ricerche (CNR) and the Canadian Natural Sciences and Engineering Research Council under grants OPG0038816 and OPG0039682. Their support is gratefully acknowledged. The computational experiments were executed at the Laboratory of Operations Research of the University of Bologna (LabOR).

References

[1] M.A. Boschetti and A. Mingozzi. The two-dimensional finite bin packing problem. Part I: New lower bounds for the oriented case. *4OR*, 1:27–42, 2003.

[2] M.A. Boschetti and A. Mingozzi. The two-dimensional finite bin packing problem. Part II: New lower and upper bounds. *4OR*, 1:135–147, 2003.

[3] N. Christofides and C. Whitlock. An algorithm for two-dimensional cutting problems. *Operations Research*, 25:30–44, 1977.

[4] G. Clarke and J.W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12:568–581, 1964.

[5] J.-F. Cordeau, M. Gendreau, A. Hertz, G. Laporte, and J.-S. Sormany. New heuristics for the vehicle routing problem. In A. Langevin and D. Riopel, editors, *Logistics Systems: Design and Optimization*. Kluwer, Boston, 2005 (to appear).

[6] J.-F. Cordeau, M. Gendreau, and G. Laporte. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30:105–119, 1997.

[7] J.-F. Cordeau and G. Laporte. Tabu search heuristics for the vehicle routing problem. In C. Rego and B. Alidaee, editors, *Metaheuristic Optimization via Memory and Evolution: Tabu Search and Scatter Search*, pages 145–163. Kluwer, Boston, 2004.

[8] J.-F. Cordeau, G. Laporte, and A. Mercier. A unified tabu search heuristic for vehicle routing problems with time windows. *Journal of the Operational Research Society*, 52:928–936, 2001.

[9] G.B. Dantzig and J.H. Ramser. The truck dispatching problem. *Management Science*, 6:80, 1959.

[10] K.A. Dowsland. Some experiments with simulated annealing techniques for packing problems. *European Journal of Operational Research*, 68:389–399, 1993.

[11] S.P. Fekete and J. Schepers. On more-dimensional packing III: Exact algorithms. *Operations Research*, 2003 (to appear).

[12] M. Gendreau, A. Hertz, and Laporte G. A tabu search heuristic for the vehicle routing problem. *Management Science*, 40:1276–1290, 1994.

[13] M. Gendreau, A. Hertz, and G. Laporte. New insertion and postoptimization procedures for the traveling salesman problem. *Operations Research*, 40:1086–1094, 1992.

[14] F. Glover. Future paths for integer programming and links to artificial intelligence. *Computers and Operations Research*, 13:533–549, 1986.

[15] F. Glover. Tabu seach – Part I. *ORSA Journal on Computing*, 1:190–206, 1989.

[16] F. Glover. Tabu seach and adaptive memory programming – advances, applications and challenges. In R.S. Barr, R.V. Helagson, and J.L. Kennington, editors, *Interfaces in Computer Science and Operations Research*, pages 1–75. Kluwer, Boston, 1996.

[17] F. Glover and M. Laguna. *Tabu Search*. Kluwer, Boston, 1997.

[18] M. Iori. *Metahueristic algorithms for combinatorial optimization problems*. PhD thesis, DEIS, University of Bologna, Italy, 2004.

[19] M. Iori, S. Martello, and M. Monaci. Metaheuristic algorithms for the strip packing problem. In P. Pardalos and V. Korotkich, editors, *Optimization and Industry: New Frontiers*, pages 159–179. Kluwer, Boston, 2003.

[20] M. Iori, J.J. Salazar González, and D. Vigo. An exact approach for the symmetric capacitated vehicle routing problem with two dimensional loading constraints. Technical Report OR/03/04, DEIS, Università di Bologna, 2003.

[21] A. Lodi, S. Martello, and D. Vigo. Heuristic and metaheuristic approaches for a class of two-dimensional bin packing problems. *INFORMS Journal on Computing*, 11:345–357, 1999.

[22] S. Martello, M. Monaci, and D. Vigo. An exact approach to the strip packing problem. *INFORMS Journal on Computing*, 15:310–319, 2003.

[23] S. Martello, D. Pisinger, and D. Vigo. The three-dimensional bin packing problem. *Operations Research*, 48:256–267, 2000.

[24] S. Martello and D. Vigo. Exact solution of the two-dimensional finite bin packing problem. *Management Science*, 44:388–399, 1998.

[25] D. Pisinger and M. Sigurd. Using decomposition techniques and constraint programming for solving the two-dimensional bin packing problem. Technical Report 03/1, Department of Computer Science, University of Copenhagen, 2003.

[26] É. D. Taillard. Parallel iterative search methods for vehicle routing problems. *Networks*, 23:661–673, 1993.

[27] P. Toth and D. Vigo. *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, 2002.